

Clock Quantum Monte Carlo: an imaginary-time method for real-time quantum dynamics

Jarrod R. McClean¹ and Alán Aspuru-Guzik¹

¹*Department of Chemistry and Chemical Biology, Harvard University, Cambridge MA, 02138*

In quantum information theory, there is an explicit mapping between general unitary dynamics and Hermitian ground state eigenvalue problems known as the Feynman-Kitaev Clock. A prominent family of methods for the study of quantum ground states are quantum Monte Carlo methods, and recently the full configuration interaction quantum Monte Carlo (FCIQMC) method has demonstrated great promise for practical systems. We combine the Feynman-Kitaev Clock with FCIQMC to formulate a new technique for the study of quantum dynamics problems. Numerical examples using quantum circuits are provided as well as a technique to further mitigate the sign problem through time-dependent basis rotations. Moreover, this method allows one to combine the parallelism of Monte Carlo techniques with the locality of time to yield an effective parallel-in-time simulation technique.

I. INTRODUCTION

Understanding the evolution of quantum systems is a central problem in physics and the design of emerging quantum technologies. However, exact simulations of quantum dynamics suffer from the so-called curse of dimensionality [1]. That is, the dimension of the Hilbert space grows exponentially with the size of the physical system. An effective remedy for the curse of dimensionality in some classical systems has been the use of Monte Carlo methods, which in many cases has an error with respect to number of samples that is independent of the dimension of the simulated system [2]. Unfortunately this favorable scaling is often lost in quantum systems of interest due to the emergence of the famous sign problem. In particular, it has hindered the use of Monte Carlo methods for fermionic systems, where it is sometimes called “the fermion sign problem”, and for real-time dynamics of general quantum systems, where it is known as “the dynamical sign problem”. The generic sign problem has been proven to belong to the computational complexity class NP-Complete [3], and recent studies of complexity have refined knowledge about the computational power of sign-problem free (or “stoquastic”) Hamiltonians [4, 5]. However, these results do not preclude the effective use of these methods on many interesting instances of physical problems.

In particular, despite the generic challenges of the sign problem, Monte Carlo methods have been used with great success in the study of electronic systems, providing a standard of accuracy in quantum chemistry and condensed matter [6–10]. In some of these methods, such as fixed node diffusion Monte Carlo, the use of a trial wavefunction allows one to approximately remove the complications of the sign problem at the cost of a small bias in the resulting energy. One alternative to such an approximation is the use of interacting walker methods [11], which attempt to solve the problem exactly without the bias introduced by a trial function. Recently, Booth et. al introduced an interacting walker method in the discrete basis of Slater determinants called Full Configuration In-

teraction Quantum Monte Carlo (FCIQMC) [12]. The sign problem in the context of this algorithm has been studied in some detail [13–15] and it has been successfully applied to both small molecular systems of chemical interest and extended bulk systems [16, 17].

The use of Monte Carlo methods to study the real-time dynamics of generic quantum systems has been comparatively less prevalent [18]. The dynamical sign problem may become more severe both with the size of the system, and duration for which it is simulated [19–21]. Despite these challenges, advances are being made in the treatment of these problems, including hybridization of Monte Carlo techniques with other methods [22–26].

The sign problem has been studied in the context of quantum computation, where it is known that a sufficient condition for efficient probabilistic classical simulation of the adiabatic evolution of a quantum system using Monte Carlo methods is that the Hamiltonian governing the quantum system is sign problem free (also known as stoquastic) and frustration free [4, 5, 27]. Projector Monte Carlo algorithms have been developed specifically for this type of problem [4, 28]. Moreover, the use of tools from quantum information allows any generic unitary evolution of a quantum system to be written as the ground state eigenvalue problem of a Hermitian Hamiltonian [29–31]. In this work, we exploit this equivalence to adapt the interacting walker method introduced by Booth et. al [12] to treat the dynamical sign problem with a method designed for the fermion sign problem.

The paper is organized as follows. First, we review the time-embedded discrete variational principle [31] and derive from it the Clock Hamiltonian [29–31], which are the essential tools for writing a general unitary evolution as a ground state eigenvalue problem of a Hermitian Hamiltonian. We then review the FCIQMC method and adapt it for application to the Clock Hamiltonian. A discussion of the theoretical and practical manifestation of the dynamical sign problem in this setting follows with numerical examples from quantum computation. Finally, we introduce a general framework of basis rotations which can be used to ameliorate the sign problem and study the

performance of this method when used in parallel computation.

II. DYNAMICS AS A GROUND STATE PROBLEM

It has been shown that any unitary quantum evolution may be formulated as a ground state eigenvalue problem with applications to classical simulation of quantum systems [31]. We briefly review the relevant results here so that this work remains self-contained.

Consider a quantum system that is described at discrete time steps t by a normalized wavefunction $|\Psi_t\rangle$. The dynamics of this system are described by a sequence of unitary operators $\{U_t\}$ such that $U_t |\Psi_t\rangle = |\Psi_{t+1}\rangle$ and $U_t^\dagger |\Psi_{t+1}\rangle = |\Psi_t\rangle$. From the properties of unitary evolution, the following is clear:

$$2 - \langle \Psi_{t+1} | U_t | \Psi_t \rangle - \langle \Psi_t | U_t^\dagger | \Psi_{t+1} \rangle = 0. \quad (1)$$

Moreover, if the wavefunctions at each point in time are only approximately known (but normalized) then

$$\sum_t \left(2 - \langle \Psi_{t+1} | U_t | \Psi_t \rangle - \langle \Psi_t | U_t^\dagger | \Psi_{t+1} \rangle \right) \geq 0 \quad (2)$$

where equality holds only in the case where the wavefunction represents an exact, valid evolution of the quantum system. To consider all moments in time simultaneously, we extend the physical Hilbert space with an ancillary quantum system to denote time. This ancillary time register takes on integer values t and is orthonormal such that $\langle t' | t \rangle = \delta_{t,t'}$. With this construction, we see that by defining

$$\mathcal{H}' = \frac{1}{2} \left(\sum_t I \otimes |t\rangle \langle t| - U_t \otimes |t+1\rangle \langle t| - U_t^\dagger \otimes |t\rangle \langle t+1| + I \otimes |t+1\rangle \langle t+1| \right) \quad (3)$$

which acts on the composite system-time Hilbert space, all valid time evolutions minimize

$$\mathcal{S} = \sum_{t,t'} \langle t' | \langle \Psi_{t'} | \mathcal{H}' | \Psi_t \rangle | t \rangle. \quad (4)$$

Note that we have adopted the convention of script letters for operators acting in the system-time Hilbert space such as \mathcal{H}' as opposed to operators only acting on the system such as U_t . The time-embedded discrete variational principle immediately follows, which simply states that this quantity is stationary under variations of the wavefunction $\delta |\Psi_t\rangle$ for all valid time evolutions, or

$$\delta \mathcal{S} = \delta \sum_{t,t'} \langle t' | \langle \Psi_{t'} | \mathcal{H}' | \Psi_t \rangle | t \rangle = 0 \quad (5)$$

To select a particular evolution of interest, one may introduce a penalty operator that fixes the state of the

system at a given time. Typically, this might represent a known initial state, and this operator in the system-time Hilbert space is given by

$$\mathcal{C}_0 = (I - |\Psi_0\rangle \langle \Psi_0|) \otimes |0\rangle \langle 0|. \quad (6)$$

The minimization of a Hermitian quadratic form constrained to have unit norm is equivalent to the eigenvalue problem for the corresponding Hamiltonian. We introduce the Lagrange multiplier λ to enforce normalization. As both \mathcal{S} and \mathcal{C}_0 are Hermitian by construction, minimization of

$$\mathcal{L} = \sum_{t,t'} \langle t' | \langle \Psi_{t'} | \mathcal{H}' + \mathcal{C}_0 | \Psi_t \rangle | t \rangle - \lambda \left(\sum_{t,t'} \langle t' | \langle \Psi_{t'} | \Psi_t \rangle | t \rangle - 1 \right) \quad (7)$$

is equivalent to solving for the eigenvector corresponding to the smallest eigenvalue of the Hermitian operator

$$\mathcal{H} = \mathcal{H}' + \mathcal{C}_0 \quad (8)$$

which we refer to as the Clock Hamiltonian. This Hamiltonian has a unique ground state with eigenvalue 0 given by the history state,

$$|\Phi\rangle = \frac{1}{\sqrt{T}} \sum_t |\Psi_t\rangle \otimes |t\rangle \quad (9)$$

which encodes the entire evolution of the quantum system. Thus, the quantum dynamics of the physical system can be obtained by finding the ground state eigenvector of \mathcal{H} .

III. FCIQMC FOR THE CLOCK HAMILTONIAN

The Full Configuration Interaction Quantum Monte Carlo (FCIQMC) method was introduced by Booth et al as a method to accurately find the ground state for electronic structure problems in a basis of Slater determinants without appealing to the fixed node approximation to eliminate the fermion sign problem [12]. We review the basics of the theory behind this method and show how it can be adapted for the Clock Hamiltonian \mathcal{H} , such that it simulates the full time evolution of a quantum system.

Let $|\Phi_i\rangle$ and λ_i be the eigenvectors and corresponding eigenvalues of \mathcal{H} . Any vector $|\Psi\rangle$ in the system-time Hilbert space acted upon by \mathcal{H} can be decomposed in terms of the eigenvectors of \mathcal{H} such that

$$|\Psi\rangle = \sum_i c_i |\Phi_i\rangle \quad (10)$$

It follows that for any $|\Psi\rangle$ not orthogonal to the ground state of the Clock Hamiltonian, $|\Phi_0\rangle$, that

$$\lim_{\tau \rightarrow \infty} e^{-\tau \mathcal{H}} |\Psi\rangle = \lim_{\tau \rightarrow \infty} \sum_i e^{-\tau \lambda_i} c_i |\Phi_i\rangle \propto |\Phi_0\rangle \quad (11)$$

Because \mathcal{H} trivially commutes with itself, we may break this operator into the successive application of many operators, such that for a large number of slices N of a finite τ ,

$$e^{-\tau\mathcal{H}} = (e^{-\frac{\tau}{N}\mathcal{H}})^N \approx (1 - \delta\tau\mathcal{H})^N \quad (12)$$

where $\delta\tau = \tau/N$. Note that the linearized time propagator used here is both simple to implement for discrete systems as well as unbiased in the final ($\tau \rightarrow \infty$) result given some restrictions on $\delta\tau$ [32]. Thus with a prescription to stochastically apply the operator

$$\mathcal{G} = (1 - \delta\tau\mathcal{H}) \quad (13)$$

repeatedly to any vector in the system-time Hilbert space, we can simulate the quantum dynamics of the physical system. τ is sometimes interpreted as imaginary-time by analogy to the Wick-rotated time-dependent Schrödinger equation, however we will only refer to τ as “simulation time” here, to avoid confusion with the simultaneous presence of real and imaginary-time.

To represent a vector in the system-time Hilbert space, we introduce discrete walkers represented by $\{i, t\}$ with associated real and imaginary integer weights $\mathcal{R}(\{i, t\})$ and $\mathcal{I}(\{i, t\})$. These walkers correspond to a single system-time configuration. The indices correspond to a system state $|i\rangle$ at time t with a complex integer weight defined by its real and imaginary components, $W(\{k, t\}) = \mathcal{R}(\{k, t\}) + i\mathcal{I}(\{k, t\})$. Given a collection set of these walkers, the corresponding normalized vector is given by

$$|\Psi\rangle = \frac{1}{K} \sum_{\{i, t\}} W(\{i, t\}) |i\rangle \otimes |t\rangle \quad (14)$$

where K is the normalization constant given by the sum of the absolute value of all the complex integer walker weights

$$K = \sum_{\{i, t\}} |W(\{i, t\})|. \quad (15)$$

We also use this notation to define matrix elements for an operator \mathcal{O} between a state $|i\rangle |t\rangle$ and $|j\rangle |t'\rangle$ as

$$\mathcal{O}_{\{j, t'\}, \{i, t\}} = \langle j | \langle t' | \mathcal{O} | i \rangle | t \rangle \quad (16)$$

To stochastically apply the operator \mathcal{G} to a vector represented by a set of such walkers, the following four steps are used, adapted from the original implementation by Booth et. al:

1. Spawning: This step addresses the off-diagonal elements of \mathcal{G} . For each walker $\{i, t\}$, we consider $N_r = \mathcal{R}(\{i, t\})$ real parents and $N_i = \mathcal{I}(\{i, t\})$ imaginary parents, both with the correct associated sign. For each of the real parents N_i , we select a coupled state at an adjacent time and attempt to

spawn a real child and imaginary child $\{j, t'\}$ with probabilities

$$p_s^{\mathcal{R}}(\{j, t'\} | \{i, t\}) = \frac{\delta\tau |\mathcal{R}(\mathcal{H}_{\{j, t'\}, \{i, t\}})|}{p_{g_t}(t', t) p_{g_s}(\{j, t'\} | \{i, t\})} \quad (17)$$

$$p_s^{\mathcal{I}}(\{j, t'\} | \{i, t\}) = \frac{\delta\tau |\mathcal{I}(\mathcal{H}_{\{j, t'\}, \{i, t\}})|}{p_{g_t}(t', t) p_{g_s}(\{j, t'\} | \{i, t\})} \quad (18)$$

with corresponding signs

$$S^{\mathcal{R}} = -\text{sign}(\mathcal{R}(\mathcal{H}_{\{j, t'\}, \{i, t\}})) \quad (19)$$

$$S^{\mathcal{I}} = -\text{sign}(\mathcal{I}(\mathcal{H}_{\{j, t'\}, \{i, t\}})) \quad (20)$$

and for each of the imaginary parents N_i we select a coupled state at an adjacent time and attempt to spawn a real child and imaginary child $\{j, t'\}$ with probabilities

$$p_s^{\mathcal{R}}(\{j, t'\} | \{i, t\}) = \frac{\delta\tau |\mathcal{I}(\mathcal{H}_{\{j, t'\}, \{i, t\}})|}{p_{g_t}(t', t) p_{g_s}(\{j, t'\} | \{i, t\})} \quad (21)$$

$$p_s^{\mathcal{I}}(\{j, t'\} | \{i, t\}) = \frac{\delta\tau |\mathcal{R}(\mathcal{H}_{\{j, t'\}, \{i, t\}})|}{p_{g_t}(t', t) p_{g_s}(\{j, t'\} | \{i, t\})} \quad (22)$$

with corresponding signs

$$S^{\mathcal{R}} = \text{sign}(\mathcal{I}(\mathcal{H}_{\{j, t'\}, \{i, t\}})) \quad (23)$$

$$S^{\mathcal{I}} = -\text{sign}(\mathcal{R}(\mathcal{H}_{\{j, t'\}, \{i, t\}})) \quad (24)$$

where probabilities $p_s > 1$ are handling by deterministically spawning $\lfloor p_s \rfloor$ walkers and spawning an additional walker with probability $p_s - \lfloor p_s \rfloor$. $\delta\tau$ is the simulation time step and may be used to control the rate of walker spawning. The functions $p_{g_t}(t', t)$ and $p_{g_s}(\{j, t'\} | \{i, t\})$ are the probability of suggesting a walker at the new time t' and of the particular state j respectively. For the Clock, an efficient choice of the time generation function, $p_{g_t}(t', t)$ is $t' = t \pm 1$ with equal probability unless the walker is at a time boundary, in which case it should move inward with unit probability. The state generation probability $p_{g_s}(\{j, t'\} | \{i, t\})$ should be chosen based on knowledge of the structure of U_t such that connected states may reach each other. In this work we use a uniform distribution where states connected by U_t are selected randomly with equal probability, however this can be refined using knowledge of U_t .

In this case, where $\{j, t'\} \neq \{i, t\}$, the matrix elements $\mathcal{H}_{\{j, t'\}, \{i, t\}}$ may be written more explicitly as

$$\mathcal{H}_{\{j, t'\}, \{i, t\}} = \begin{cases} -\frac{1}{2} \langle j | U_t | i \rangle & t' = t + 1 \\ -\frac{1}{2} \langle j | U_t^\dagger | i \rangle & t' = t - 1 \\ 0 & \text{otherwise} \end{cases} \quad (25)$$

2. Diagonal Death/Cloning: This step addresses the application of the diagonal of \mathcal{G} . In this step, for each parent walker $\{i, t\}$ (not yet including child walkers spawned in the last step), calculate

$$p_d(\{i, t\}) = \delta\tau(\mathcal{H}_{\{i, t\}, \{i, t\}} - S) \quad (26)$$

where S is a shift that is used to control the population in the simulation. Now for each real and imaginary parent N_r and N_i associated with $\{i, t\}$, if $p_d > 0$, the parent is destroyed. If $p_d < 0$, the parent is cloned with a probability $|p_d|$, handling instances of greater than unit probabilities as in the previous step.

In the case of the Clock, the diagonal matrix elements take on a simple form

$$\mathcal{H}_{\{i, t\}, \{i, t\}} = \begin{cases} 1/2 + (1 - |\langle i | \Psi_0 \rangle|^2) & t = 0 \\ 1/2 & t = T - 1 \\ 1 & \text{otherwise} \end{cases} \quad (27)$$

3. Annihilation: In this step, all previously existing and newly spawned walkers are searched, and any which match are combined such that both their real and imaginary components are summed. In the event that any walker ends up with 0 total weight, it is removed entirely from the simulation. In the case of the Clock, it is advantageous to store walkers grouped by time, such that in parallel implementations the simulation can be easily split along this dimension. This will be elaborated upon later. Within each group it is advisable to use any natural ordering present on the basis states to enable binary search that can locate identical walkers in a time that is logarithmic in the number of walkers at a given time. Alternatively one can use hash tables to facilitate annihilation [33].

A single iteration of the above algorithm is cartooned in Fig. 1. By using this procedure, the operator G is iteratively applied until the state of walkers is equilibrated at some simulation time $\tau > \tau_{eq}$, with a number of walkers N_w . The average of some observable O may be estimated at simulation time τ according to

$$\langle O \rangle(\tau) = \frac{\langle \Phi(\tau) | O | \Phi(\tau) \rangle}{\langle \Phi(\tau) | \Phi(\tau) \rangle} \quad (28)$$

By averaging over the simulation time τ and correcting for the autocorrelation time of the quantity $\langle O \rangle$ using standard statistical procedures, the average may be converged to the desired precision. In general, however, the simulation time averaged quantity $\langle O \rangle_\tau$ may be biased due to the sign problem [13–15]. This bias may be removed to an arbitrary degree by increasing the number of walkers N_w such that the state remains sign-coherent between steps. The number of walkers required to remove the bias to a given precision depends both on the severity of the sign problem and the amount of Hilbert space




FIG. 1. A schematic representation of a single iteration of the FCIQMC algorithm for the Clock. The larger squares represent real-time, and sub-squares represent the possible quantum states occupied by either positive (blue) or negative (red) walkers. In each iteration, the set of parents spawns potential children to adjacent times, with parentage being indicated by dotted lines. Simultaneously the set of parent walkers are considered for diagonal death. Finally, the remaining set of parents and spawned children are combined, allowing walkers with opposing signs at the same state and time to annihilate.

the physical problem occupies [13–15]. To this end, we define a problem-dependent number n_c such that when $N_w > n_c$, the time averaged quantity $\langle O \rangle_\tau$ is accurate to the desired precision. Because this is an NP-Complete problem, one must expect that in general, n_c is on the order of the dimension of the Hilbert space, D , that is, it grows exponentially with the size of the system and linearly with real-time. However, for many systems of interest in ground state problems it has been found that $n_c \ll D$ [15–17], and one might expect the same to be true for some dynamical problems. We now turn our attention to the scaling and properties of n_c for dynamical systems.

IV. MANIFESTATION OF THE SIGN PROBLEM

The conditions for the efficient simulation of a Hamiltonian on a classical computer have been studied in the context of quantum complexity theory. It is known that if a Hamiltonian is frustration free and has real, non-positive off-diagonal elements in a standard basis (stoquastic) that it may be probabilistically simulated to a set precision in a time that is polynomial in the size of

the system [4, 5].

For practical purposes, there are limitations on the system operators one may simulate. In particular, the system operators must be the sum of a polynomial number of terms. This simply originates from the need to be able to efficiently evaluate matrix elements of a given state. The interaction of at most k particles, or k -local interactions, in the physical system is a sufficient but not necessary condition for this to be true. The Clock construction has also been recently generalized to open quantum systems [34], where even in this case a 2-local construction is generally possible with the use of gadgets. Alternatively, if the Clock is constructed from a sequence of unitary gates that act on at most k qubits in quantum computation, then the Clock will also satisfy this requirement.

The presence of frustration in interacting systems can cause the autocorrelation time of measured observables to diverge exponentially, rendering their efficient simulation intractable even in cases where the Hamiltonian is bosonic or sign problem free [3, 35]. It has been proven generally that the Clock Hamiltonian is frustration free, with a unique ground state separated from the first excited state with a gap of $O(1/T^2)$ where T is the number of discrete time steps being considered at once.

If an operator is stoquastic (or sign problem free), then the off-diagonal elements that correspond to transitions in a Monte Carlo simulation all be non-positive. The operator \mathcal{G} will contain only positive transition probabilities in this case and have a ground state corresponding to a classical probability distribution by the Perron-Frobenius theorem [4, 5]. In the context of the FCIQMC method introduced, this means that walkers will never change signs throughout the simulation, and all averages will be sign-coherent and unbiased independent of the number of walkers N_w . In the Clock Hamiltonian, the off-diagonal elements correspond to the set of unitary operators with their adjoints $\{U_t, U_t^\dagger\}$, and the penalty term \mathcal{C}_0 . For the standard computational initial state ($|0\rangle^{\otimes N}$), the penalty term \mathcal{C}_0 has a fixed sign, and thus the Clock Hamiltonian is stoquastic if $\{U_t, U_t^\dagger\}$ represented in the standard basis has all real positive entries, yielding non-positive off-diagonal entries in the Clock.

Given the ubiquity of k -local interactions in physical problems and the rigorous proof that the Clock Hamiltonian is frustration free, we will take these two conditions as given and consider more carefully the stoquastic condition. Consider a simple example of a unitary evolution that may be simulated on a classical computer efficiently, namely reversible classical computational. All reversible classical circuits may be expressed in terms of Toffoli gate sequences, which is unitary and acts to switch a target bit conditional on the state of two control bits. In the standard computational basis it has a representation given

by

$$Toff = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix} \quad (29)$$

The Clock when constructed with unitary Toffoli gates is stoquastic and $n_c \approx 1$. Although a stoquastic Hamiltonian is sufficient for this to be the case, it is not a necessary condition. To see this, consider a slightly different set of unitary operators, namely the standard Pauli group gates, X_i , Y_i , and Z_i in combination with the CNOT gate. These gates have the following unitary representations in the standard computational basis

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad (30)$$

$$Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \quad (31)$$

$$Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \quad (32)$$

$$CNOT = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} \quad (33)$$

Considering for now only the computational basis we simulate in (a restriction we later lift), it is clear that given the complex entries and varying signs of the off-diagonals, that a Clock Hamiltonian built from this gate set will not be stoquastic if even a single Y or Z gate is used. However these gates also have the property that they map single configurations to single configurations, and as a result no interference occurs, yielding all sign coherent averages and $n_c \approx 1$. We call this type of transformation, which is configuration number preserving, “quasi-classical”, in contrast to classical which we define as configuration number preserving as well as phase preserving. Thus a stoquastic Clock Hamiltonian is a sufficient, but not necessary condition for the simulation procedure to be sign-problem free.

Consider a slightly more general local rotation R parameterized by an angle θ

$$R(\theta) = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \quad (34)$$

In this case, the value of n_c as a function of system size is more complex. These represent the real-time evolutions of local spin Hamiltonians for systems of spin- $\frac{1}{2}$ particles.

In Fig. 2 we consider a single rotation $R(\theta)$ with $\theta = 5\pi/32$ applied uniformly to 11 qubits initialized to $|0\rangle^{\otimes N}$. As the Clock Hamiltonian in this simulation is




FIG. 2. Computed expectation value for S_z for a single qubit at the final time in the simulation as a function of the average number of walkers kept in the simulation. There are 11 total qubits in the simulation. It is apparent the system exhibits a sharp transition between a totally sign incoherent sampling where all averages become zero, and a sign coherent region where the averages begin to converge to the appropriate value.

neither stoquastic nor quasi-classical, one observes a sign-incoherent region for a small number of walkers, where all averages tend towards 0, until some critical threshold $N_w > n_c$ is reached where a transition occurs to sign-coherent sampling, and the average converges to the true value. We note that some implementations of the FCIQMC algorithm have used the diagonal shift S as a proxy for convergence [12], but we did not observe a similar plateau trend here. The history state being sampled in this case is given by

$$|\Psi\rangle = \sum_t \frac{1}{\sqrt{T}} (R(\theta) |0\rangle)^{\otimes t} |0\rangle^{\otimes T-t} \quad (35)$$

The formal structure of this evolution is quite similar for all values of θ , however the states that result can exhibit quite different features with respect to the sign problem in sampling. In Fig. 3 we examine the same circuit, but include many different rotation angles. One sees that not only does the value n_c change as a function of rotation angle, but the rate of the transition is quite different as well, favoring sharper, earlier transitions for rotations that are closer to quasi-classical ($\theta = 0$).

V. MITIGATING THE SIGN PROBLEM

In the last section we considered the impact of sign problem as it related to local rotations (or the dynamics of distinguishable non-interacting particles). The apparent challenges in this domain are unsettling given that trivial solutions are known for this problem. Here we




FIG. 3. Computed expectation value for S_z for a single qubit at the final time in the simulation as a function of the average number of walkers kept in the simulation and the rotation angle used in the simulation. The rotation angle θ is indicated by the line label. The simulation contains 11 total qubits for all rotation angles. One sees that the closer the rotation is to quasi-classical, the sharper and earlier the transition to sign coherent sampling.

propose a simple scheme to mitigate the sign problem to an arbitrary extent using preliminary computation.

It is known that the sign problem is generically basis dependent. To this end we propose an analogous approach to the interaction picture in quantum dynamics, where the walkers at each point in time are expressed in a new basis determined by a generic time-dependent unitary rotation given by $\{\mathcal{B}_t\}$. The evolution operators are also dressed by this change such that in the new basis, the Clock is constructed from the rotated operators given by

$$U'_t = \mathcal{B}_{t+1}^\dagger U_t \mathcal{B}_t \quad (36)$$

Moreover, the computation of any Hermitian observable O must also take into consideration the new basis, such that

$$\langle O \rangle(\tau) = \frac{\langle \Phi(\tau) | \mathcal{B}_t^\dagger O \mathcal{B}_t | \Phi(\tau) \rangle}{\langle \Phi(\tau) | \Phi(\tau) \rangle} \quad (37)$$

If one finds a set of $\{\mathcal{B}_t\}$ that renders the Clock Hamiltonian stoquastic or quasi-classical, the resulting Hamiltonian may be sampled readily. One expects that in general, finding this basis must be at least as difficult as solving exactly the problem of the quantum evolution. In fact, it is easy to see that one may choose the exact evolution as the set of basis rotations, and that this renders the Clock Hamiltonian stoquastic and trivial, such that the evolution is dictated by the identity at all times. Of course the price one must pay for this is that the computation of observables requires the full evolution to be known.




FIG. 4. The quantum circuit diagram for the circuit used to test the efficacy of time-dependent local rotations in ameliorating the sign problem. The angle used in this case is $\theta = 0.49$. We compare the results from this circuit as a function of the number of controls that are removed from the NOT gates (crossed circle here), and whether time dependent local basis rotations are utilized. The controls are removed from the end of the circuit first.

However, as was seen above, it is not necessary for the Clock to be rendered completely trivial. Even approaching a quasi-classical Hamiltonian in an approximate sense can greatly reduce the sampling costs. For some instances, one may find an approximate set of rotations that make the Clock nearly stoquastic or quasi-classical, and the remainder of the sign problem can be handled by maintaining a reasonable number of walkers N_w in the simulation. As an example of this procedure, we consider the simple case where $\{\mathcal{B}_t\}$ are determined entirely by the local rotations in a quantum circuit. Specifically, for local rotations, $\mathcal{B}_t = \prod_{t' < t}^0 U_{t'}$, where $U_{t'}$ is replaced by I for two- or more qubit operations. It is clear that for circuits consisting of only local rotations, as in the previous section, this is equivalent to exact evolution and the resulting Clock Hamiltonian becomes trivial ($U_t = I \forall t$). To study how this works in the non-trivial case, we examine a similar circuit of local rotations, but now with a variable number of CNOT gates included. This elucidates to what extent the use of basis rotations can mitigate the sign problem when they are not an exact solution to the dynamics considered. A depiction of this circuit is given in Fig. 4.

In Fig. 5 we consider a simple quantum circuit consisting of a series of local rotations followed by NOT gates, with a variable number of the NOT gates under control. With the given rotation angle ($\theta = 0.49$), these are entangling operations not covered by the simple local basis rotation scheme we use here. However, it is seen that even for 8 controlled NOT gates, the use of local basis rotations dramatically reduces the number of walkers required to reach sign-coherent sampling, indicating this scheme can be computationally effective even for simulations containing a considerable fraction of two-qubit entangling operations. Further tests of more complex quantum circuits are needed to determine the efficiency of different rotation schemes as a function of the structure of the quantum circuit.




FIG. 5. The mean value of a spin observable is plotted as a function of the mean number of walkers labeled by the number of CNOT gates both with local basis rotation(-Rot) and without. It is seen that even for a relatively high proportion of CNOT gates, the rotated basis performs far better than the non-rotated basis, requiring a lower number of walkers to reach sign-coherent sampling.

VI. PARALLEL-IN-TIME SCALING

Monte Carlo methods are often championed as the ultimate parallel algorithms, associated with the phrase “embarrassingly parallel”. Given the evolution of modern computational architectures towards many-core architectures with slower Clock speeds, Monte Carlo will continue to play a growing role in the numerical simulation of physics at the boundaries of our computational capabilities. Interacting walker Monte Carlo methods, can be more difficult to parallelize effectively due to the annihilation step where communication of walkers is unavoidable.

In contrast to the most general interacting walker algorithm, which may require heavy communication between all processes, the FCIQMC method applied to the Clock Hamiltonian may take advantage of time-locality to create an efficient parallel-in-time algorithm using the standard method of domain decomposition in time. Using this construction, only processes containing adjacent times need to communicate their child walkers, which may be done simultaneously in a time that is constant for the number of processes involved. This remains true so long as the number of time steps under consideration is larger than the number of processes in use, which is typically the case. In the case that the number of processes is much greater than the number of timesteps, this scheme may still be used by blocking multiple processors to each time, and utilizing an all-to-all communication pattern within each block only. The difference between these two communication patterns is highlighted in Fig.

6.




FIG. 6. A schematic representation of the communication patterns the annihilation step of interacting walker Monte Carlo schemes, where the boxes represent different MPI processes and the ellipsis represents the rest of the processes. In the case of the Clock (a), a time domain decomposition allows one to restrict communication to only nearest neighbour processes, facilitating simple, constant time communication amenable to the architecture of modern parallel computers. In the more general case (b), a clear partitioning may not be readily achievable, and all processes may need to communicate with all other processes, creating a bottleneck.

To demonstrate the scaling properties of this approach, we consider the scaling as a function of the number of processors for fixed total problem size, or strong scaling, with our implementation. This benchmarking is done on a standard Linux cluster composed of AMD Opteron 6376 processors. The parallel speedup with respect to single core time as a function of the number of processors is given in Fig 7. Here, we see that we are able to combine the parallelism of Monte Carlo with the locality of time-decomposition to achieve practical parallel efficiencies of over 95% with 2 processors and 70% with 8 using a simple MPI implementation on a commodity cluster.

VII. CONCLUSIONS

In this work we reviewed the mapping between unitary dynamics and ground state eigenvalue problems. We then showed how the FCIQMC method, a technique originally designed to ameliorate the fermionic sign problem for ground state electronic systems, could be applied to quantum dynamics problems as a direct result of this mapping. This establishes a potential research direction




FIG. 7. A scaling study of our method implementation with a fixed total problem size (strong scaling), showing parallel efficiencies and speedups. The simulation consisted of 11 qubits with 128 time points generated by consecutive local rotations with $\theta \approx 0.098$. The simulation maintained on average 10^6 walkers in each simulation-time step and the wall clock time was measured to the point of an equivalent number of statistical samples.

for explicit connections between the fermionic and dynamical sign problems that plague quantum Monte Carlo simulations, and provides a pathway for the transfer of tools between the two domains.

The numerical consequences of the dynamical sign problem in this context were studied using a few basic quantum circuits. It was found that even local rotations can exhibit a severe sign problem depending on the form of the rotation and how different it is from a quasi-classical operation. We then introduced a general method analogous to the interaction picture in dynamics or natural orbitals in the study of eigenstates that uses basis rotations to mitigate the difficulty of the problem. The costs and benefits of different types of rotations require further research, however we showed that even local rotations can have a significant benefit for non-trivial circuits. Finally, we discussed the structure of the problem in the context of parallel-in-time dynamics, and showed high parallel efficiencies with only a basic MPI implementation on a commodity cluster.

Overall, we believe this is a promising new method for the simulation of quantum dynamics. It clarifies the bridge between dynamics and ground state problems and is capable of effectively utilizing parallel computing architectures. While we have only demonstrated it for quantum circuits, we believe it will be generally useful for the study of quantum dynamics.

VIII. ACKNOWLEDGEMENTS

We acknowledge Ryan Babbush and Thomas Markovich for their helpful comments on the manuscript. J.M. is supported by the Department of Energy Com-

putational Science Graduate Fellowship under grant number DE-FG02-97ER25308. A.A-G. acknowledges the support of the Air Force Office of Scientific Research under Award No. FA9550-12-1-0046 and the National Science Foundation under Award No. CHE-1152291.

-
- [1] D. Tannor, *Introduction to Quantum Mechanics: A Time-Dependent Perspective* (University Science Books, 2007).
 - [2] K. Binder and D. W. Heermann, *Monte Carlo simulation in statistical physics: an introduction* (Springer, 2010).
 - [3] M. Troyer and U. J. Wiese, Phys. Rev. Lett. **94**, 170201 (2005).
 - [4] S. Bravyi, D. P. Divincenzo, R. Oliveira, and B. M. Terhal, Quantum Info. Comput. **8**, 361 (2008).
 - [5] S. Bravyi and B. Terhal, SIAM J. Comput. **39**, 1462 (2009).
 - [6] B. Hammond, W. Lester, and P. Reynolds, *Monte Carlo Methods in Ab Initio Quantum Chemistry*, World Scientific Lecture and Course Notes in Chemistry ; Vol. 1 (World Scientific, 1994).
 - [7] M. Nightingale and C. Umrigar, *Quantum Monte Carlo Methods in Physics and Chemistry*, Nato ASI series (Kluwer Academic Publishers, 1999).
 - [8] S. Baroni and S. Moroni, Phys. Rev. Lett. **82**, 4745 (1999).
 - [9] W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, Rev. Mod. Phys. **73**, 33 (2001).
 - [10] S. Zhang and H. Krakauer, Phys. Rev. Lett. **90**, 136401 (2003).
 - [11] S. Zhang and M. H. Kalos, Phys. Rev. Lett. **67**, 3074 (1991).
 - [12] G. H. Booth, A. J. W. Thom, and A. Alavi, J. Chem. Phys. **131**, 054106 (2009).
 - [13] J. S. Spencer, N. S. Blunt, and W. M. Foulkes, J. Chem. Phys. **136**, 054110 (2012).
 - [14] M. H. Kolodrubetz, J. S. Spencer, B. K. Clark, and W. M. Foulkes, J. Chem. Phys. **138**, 024110 (2013).
 - [15] J. J. Shepherd, G. E. Scuseria, and J. S. Spencer, ArXiv e-prints (2014), arXiv:1407.4800 [physics.comp-ph].
 - [16] G. H. Booth and A. Alavi, J. Chem. Phys. **132**, 174104 (2010).
 - [17] G. H. Booth, A. Gruneis, G. Kresse, and A. Alavi, Nature **493**, 365 (2013).
 - [18] C. H. Mak, Phys. Rev. Lett. **68**, 899 (1992).
 - [19] R. Feynman and A. Hibbs, *Quantum Mechanics and Path Integrals: Emended Edition* (Dover Publications, Incorporated, 2012).
 - [20] D. Thirumalai and B. J. Berne, J. Chem. Phys. **79**, 5029 (1983).
 - [21] N. Makri and W. H. Miller, Chem. Phys. Lett. **139**, 10 (1987).
 - [22] N. Makri and W. H. Miller, J. Chem. Phys. **89**, 2170 (1988).
 - [23] C. H. Mak and D. Chandler, Phys. Rev. A **41**, 5709 (1990).
 - [24] N. Makri, Comp. Phys. Comm. **63**, 389 (1991).
 - [25] N. Makri, J. Phys. Chem **97**, 2417 (1993).
 - [26] V. Jadhao and N. Makri, J. Chem. Phys. **129**, 161102 (2008).
 - [27] A. M. Childs, D. Gosset, and Z. Webb, Science **339**, 791 (2013).
 - [28] S. Bravyi, ArXiv e-prints (2014), arXiv:1402.2295 [quant-ph].
 - [29] R. Feynman, Int. J. Theor. Phys. **21**, 467 (1982).
 - [30] A. Kitaev, A. Shen, M. Vyalii, and N. Vyalii, *Classical and Quantum Computation*, Graduate Studies in Mathematics (American Mathematical Society, 2002).
 - [31] J. R. McClean, J. A. Parkhill, and A. Aspuru-Guzik, Proc. Natl. Acad. Sci. U.S.A. **110**, E3901 (2013).
 - [32] N. Trivedi and D. M. Ceperley, Phys. Rev. B **41**, 4552 (1990).
 - [33] G. H. Booth, S. D. Smart, and A. Alavi, ArXiv e-prints (2013), arXiv:1305.6981 [physics.comp-ph].
 - [34] D. G. Tempel and A. Aspuru-Guzik, ArXiv e-prints (2014), arXiv:1406.5631 [quant-ph].
 - [35] T. C. Wei, M. Mosca, and A. Nayak, Phys. Rev. Lett. **104**, 040501 (2010).